

Do Markets Believe in Transformative AI?

Isaiah Andrews and Maryam Farboodi*

Abstract

Economic theory predicts that transformative technologies may influence interest rates by changing growth expectations, increasing uncertainty about growth, or raising concerns about existential risk. Examining US bond yields around major AI model releases in 2023-4, we find economically large and statistically significant movements concentrated at longer maturities. The median and mean yield responses across releases in our sample are negative: long-term Treasury, TIPS, and corporate yields fall and remain lower for weeks. Viewed through the lens of a simple, representative agent consumption-based asset pricing model, these declines correspond to downward revisions in expected consumption growth and/or a reduction in the perceived probability of extreme outcomes such as existential risk or arrival of a post-scarcity economy. By contrast, changes in consumption growth uncertainty do not appear to drive our results.

1 Introduction

Since the debut of ChatGPT in November 2022, generative AI models have attracted intense interest from policymakers, researchers, and businesses. Some discussions of these models have raised the possibility AI could lead to an increase, perhaps even a dramatic acceleration, in the rate of economic growth (Brynjolfsson et al., 2019; Trammell and Korinek, 2023; Acemoglu and Lensman, 2024; Jones, 2024; Korinek and Suh, 2024). Other discussions have suggested the possible gains from AI may be overstated, and argued that widespread AI adoption could potentially slow economic growth (Acemoglu and Restrepo, 2020). Many authors

*This version: September 4, 2025. First version: July 14, 2025. Previous versions of this paper circulated with the title “Pricing Transformative AI.” Andrews: MIT Department of Economics and NBER; iandrews@mit.edu. Farboodi: MIT Sloan School of Business, NBER, and CEPR; farboodi@mit.edu. We thank Marcellus Andrews, Luigi Bocola, Markus Brunnermeier, Valentin Haddad, Basil Halperin, Chad Jones, Peter Kondor, Lukasz Kowalik, Pablo Kurlat, Monika Piazzesi, Stefan Nagel, Tano Santos, Jesse Shapiro, and Cheryl Smith for very helpful discussions of and comments on this project.

have even raised the possibility that poorly understood- and controlled-AI could pose an existential risk to humanity (Acemoglu and Lensman, 2024; Jones, 2024; Kokotajlo et al., 2025).

It can be unclear to what extent the enthusiasm around AI reflects a genuine belief in its transformative potential, as opposed to belief in profit opportunities that may not translate into widespread or persistent growth. While the future impacts of AI are inherently unknown, understanding the beliefs of market participants is a potentially valuable input to both policy and research discussions. Financial market data has been used to infer market beliefs in other settings (Jackwerth and Rubinstein, 1996; Wolfers and Zitzewitz, 2004; Gürkaynak et al., 2010; Binsbergen et al., 2012; Van Binsbergen et al., 2013), but empirical evidence about market beliefs on transformative AI is limited. In this paper, our goal is to use financial market data to provide systematic evidence regarding the beliefs of market participants about the possibility of transformative AI, by which we will mean AI technologies with a large and sustained impact on living standards, particularly through impacts on consumption growth or existential risk. The premise of our analysis is that if economic actors take seriously the possibility of transformative AI, this should be reflected in a wide range of forward-looking behaviors and, consequently, in long-term asset prices, including assets such as US Treasury bonds which are not directly connected to AI.

That beliefs about transformative AI should affect agents' optimal choices is pointed out by e.g. Jones (2024). Chow et al. (2024) combine this observation with classic insights from consumption-based asset pricing to relate risk-free interest rates to market beliefs about transformative AI. The intuition is simple: if agents think AI will dramatically increase the rate of economic growth, then (on average across the economy) agents must expect to be richer in the future than they are today. This should decrease the marginal value of future consumption relative to present consumption, so real interest rates must rise in equilibrium. On the other hand, if agents think AI poses an existential risk, and so doubt that they will be alive in the future, this should also drive up interest rates. Thus, both higher growth expectations and more concern for existential risk should increase real interest rates. Beyond expectations, uncertainty also matters: if AI increases agents' uncertainty about future consumption, this will fuel precautionary saving and so decrease real interest rates (see, e.g., Gil, 2024).

Motivated by these predictions, we study the behavior of US Treasury yields around major model release dates for five major AI labs (OpenAI, Anthropic, Google DeepMind, xAI, and DeepSeek) in calendar years 2023 and 2024. As shown in Figure 1, we find that US Treasury yields substantially *decline* around model release dates, with a median




Figure 1: Median change in yields relative to fifteen (trading) days before event for constant-maturity US Treasury Bonds. Median taken across 15 major AI release events for which a ± 15 trading day window is contained in calendar years 2023-4.

decline across model release dates exceeding 10 basis points, or 0.1 percentage points, for most series. These declines are economically large and persist through 15 (trading) days after the model release. We find similar results for Treasury Inflation Protected Securities (TIPS) and corporate bond yields. Yield movements appear to begin before the release of the model, which may not be surprising given that for at least some releases, we know that models were made available to outside experts prior to the release date.

Under the assumption that AI model release dates are as good as random, and in particular that they are unrelated to other factors which may influence bond yields, the yield changes shown in Figure 1 reflect causal effects of AI releases and associated information. Under the same assumption, we further show these impacts are statistically significant, particularly for longer-maturity bonds. To probe the causal interpretation of our results we conduct a range of robustness checks, including dropping subsets of model releases, comparing to alternative date series, controlling for other information that might have influenced bond yields, and repeating our analysis on sub-samples of the data. Throughout we find evidence of negative, and often statistically significant, yield responses to AI model releases.

If bond yields drop in response to AI model releases, what does this imply for investor beliefs? To answer this question, we interpret our estimates through a simple equilibrium model of asset prices. We first show that “doom” (i.e. existential risk) and “bliss” (i.e. extremely fast growth which eliminates material scarcity, and which e.g. Jones 2024 discusses as a singularity) have equivalent asset pricing implications, since both drive the marginal utility of consumption to zero. Thus, we cannot hope to disentangle beliefs about “bliss” and “doom” via asset prices. We then use a restrictive version of this model (assuming, *inter alia*, a representative agent with CRRA utility) to quantitatively interpret our empirical results. We show that, under these additional assumptions, changes in both growth expectations and the perceived probability of extreme “bliss” or “doom” outcomes lead to a level shift in a forward yield curve. By contrast, since the consumption growth impacts of AI compound over time, changes in uncertainty about AI growth effects imply changes in the forward curve slope. Applied to our data, this model suggests that the average AI model release in our sample led investors to think that (i) expected consumption growth is lower, (ii) extreme “doom” or “bliss” outcomes are less likely, or both. In particular, the model implies that the average model release led to an approximately 0.208 percentage point drop in the annual probability of “bliss” or “doom,” or a $0.208/\gamma$ percentage point decline in the expected annual rate of consumption growth for γ the CRRA coefficient of the representative agent. By contrast, we find that model releases have little to no effect on the slope of the forward curve, suggesting that changes in consumption growth uncertainty do not explain our results.

Taken together, these results suggest that investors, in aggregate, do take seriously the possibility of transformative AI, since new information about AI models has an economically and statistically significant impact on non-AI-related asset prices. A simple model suggests, however, that the primary direction of updating across the model release dates we study was towards lower consumption growth or a lower probability of “bliss” and “doom,” rather than towards greater consumption growth uncertainty.¹

To the extent investors lowered their growth expectations around the model releases in our sample this raises a natural question. Did investors think AI advances would be good for consumption growth, but find the rate of technological progress disappointing? Or were they positively surprised by the rate of progress but pessimistic about the consumption growth implications? While we do not have direct evidence on investor beliefs, using

¹This is not necessarily incompatible with rising stock prices for AI-related firms, since one could think these firms will be highly profitable for reasons which need not imply sustained growth effects. See Section 6 of Chow et al. (2024) for further discussion of why the relationship between AI expectations and equity prices may be ambiguous.

complementary data from the online forecasting platform Metaculus we show that certain AI capability forecast timelines shifted earlier around model releases in our sample, though others show no effect. This suggests that this group of observers, at least, was positively surprised by some aspects of AI progress.

Our empirical findings admit alternative interpretations. While we think a causal interpretation of our results is plausible, we cannot rule out that there is some other force behind yield changes around model release dates. Even granting a causal interpretation, there are many ways that reality deviates from our simple model, and these deviations may suggest alternative explanations for the effects we document. For instance we assume markets are complete, while in reality a number of forces such as labor market frictions or non-competitive behavior could impede risk-sharing among agents and potentially explain our findings. Any alternative explanation for the patterns in Figure 1 must, however, account for large, sustained yield declines in one of the most liquid markets in the world around AI model release dates, and so may be of interest in its own right.

Literature Review While there is little prior evidence about the aggregate impact of AI on the economy, or on market perceptions of that impact, there is a small but growing literature that uses data on job postings or asset prices to study the impact of AI on labor outcomes and compensation, as well as on firm growth (Webb, 2019; Acemoglu et al., 2022; Babina et al., 2023, 2024; Eisfeldt et al., 2024; Hampole et al., 2025).

The influence of growth prospects on financial markets is a widely discussed topic in the consumption-based asset pricing literature (see Mehra 2012 for a summary and Duffie 2010 for a textbook treatment). An important observation in this literature is that agents' discount factors, expected growth, and perceived growth uncertainty all influence prevailing interest rates. Chow et al. (2024) abstract away from consumption growth uncertainty and show that discounting (e.g. due to existential risk) and growth expectations impact interest rates the same way in the context of transformative AI. They further show that, consistent with their theoretical analysis, real interest rates are positively correlated with both growth expectations and realized growth in a cross-country analysis. Other recent work, by contrast, finds a modest or negative relationship between growth and real interest rates (Hamilton et al., 2015; Bruce and Hansen, 2013; Borio et al., 2017; Lunsford and West, 2019; Rogoff et al., 2024). Our main contribution is to document how news about AI progress impacts interest rates. We then interpret these impacts through the lens of a consumption-based asset pricing model.

Our analysis also relates to the literature studying the impact of macroeconomic announcements on financial returns. In the context of Treasury yields and using FOMC announcements, Lucca and Moench (2015) document no statistically significant pre-FOMC announcement drift for Treasury bonds in the 1994-2011 period, while Savor and Wilson (2013) provide evidence of small announcement premiums for Treasury bonds, averaging about 3 basis points on announcement days, using data from 1961-2009.

The rest of the paper is organized as follows. Section 2 introduces a simple equilibrium asset pricing model and uses it to predict the effects of transformative AI on bond yields. Section 3 describes our data and empirical strategy, including the permutation tests we use to assess statistical significance. Section 4 reports our main empirical results, while Section 5 quantitatively interprets these results in a simplified version of our model and provides evidence on AI forecast timelines. Finally, Section 6 provides additional discussion.

2 Transformative AI in a Dynamic Economy

This section lays out a simple model of a dynamic stochastic economy, and shows that this model makes stark predictions about how investor beliefs about the possibility of transformative AI translate to asset prices. As we discuss in the introduction, by *transformative AI* we mean AI technology that substantially changes the future trajectory of the economy. Specifically, following Jones (2024) and Chow et al. (2024) we consider the possibility that AI may (i) substantially change the rate of economic growth or even (ii) lead to a more radical shift, such as the extinction of humanity (“doom”) or the arrival of a post-scarcity economy (“bliss”).

To model these possible impacts from AI, following e.g. Chapter 2 of Duffie (2010), we consider a discrete-time economy over periods $t=0,1,\dots,\bar{T}$, with uncertainty described by a probability space $(\Omega,\mathcal{F},\mathbb{P})$ where the $\mathcal{F}_t \subseteq \mathcal{F}$ denotes the set of events which are known at period t . For simplicity we assume a finite number of states $\omega \in \Omega$ and agents i . To capture the possibility of “doom” and “bliss,” similar to Jones (2024) we assume each agent i has time-separable utility

$$\mathbb{E}_0 \left[\sum_{t=1}^{\bar{T}} \beta^t (1\{t \leq T\} u_i(C_{i,t}) + 1\{t > T\} U_{i,t}^*) \right],$$

where \mathbb{E}_t denotes the conditional expectation given \mathcal{F}_t , and $T \leq \bar{T}$ denotes the (random)

date after which “doom” or “bliss” occurs. We assume that u_i is increasing and concave for all i with $\lim_{c \rightarrow \infty} u'_i(c) = 0$, while flow-utilities $U_{i,t}^*$ after T are independent of asset holdings. We henceforth normalize these post- T flow utilities to zero and write agent i 's utility as $\mathbb{E}_0 \left[\sum_{t=1}^T \beta^t u_i(C_{i,t}) \right]$.

We assume complete markets and absence of arbitrage. By standard arguments (Duffie, 2010), this implies that there exists a stochastic discount factor (SDF) that prices all assets. In particular if we consider an asset that pays Y_{t+h} units of consumption in period $t+h$ and nothing at any other time, its period t price is given by

$$V_t(Y_{t+h}) = \mathbb{E}_t[M_{t,t+h} Y_{t+h}], \quad (1)$$

where $M_{t,t+h}$ is the SDF t to $t+h$, for simplicity we write $M_{t+1} \equiv M_{t,t+1}$, and $M_{t,t+h} = \prod_{s=1}^h M_{t+s}$ cumulates the one-step-ahead SDFs. More generally, let $Y = \{Y_{t,h}\}_{h=0}^{\bar{T}}$ denote a general stream of payoffs $Y_{t,h}$ for periods $h=0, \dots, \bar{T}-t$. The asset with this stream of payoffs has time- t price $V_t(Y) = \sum_{h=0}^{\bar{T}-t} V_t(Y_{t+h})$.

Standard arguments further imply that in equilibrium the SDF coincides with the marginal rate of substitution for a representative agent with utility $\mathbb{E}_0 \left[\sum_{t=1}^T \beta^t u(C_t) \right]$, where $C_t = \sum_i C_{i,t}$ is aggregate consumption and $u(C_t) = \sum_i \lambda_i u_i(C_{i,t})$ for $\lambda_i \geq 0$, where u is increasing and concave by construction. That is, we can write the SDF as

$$M_{t,t+h} = \beta^h \frac{u'(C_{t+h})}{u'(C_t)} 1_{\{t+h \leq T\}}, \quad (2)$$

so there is a direct relationship between aggregate consumption C_t , the “doom” or “bliss” date T , and the SDF.²

Equation (2) has two important implications. First, note that the extreme possibilities of “doom” and “bliss” both enter only through the date T after which asset holdings are irrelevant. Consequently, beliefs about “doom” and “bliss” have identical asset pricing implications. Hence, under this model we have no hope of telling the two apart based on asset prices. Second, note that since the representative agent's flow utility u is increasing and concave, increases in future aggregate consumption C_{t+h} lead to a decrease in the SDF.

²Indeed, this follows from the fact that each agent's marginal utility obeys the same equality,

$$M_{t,t+h} = \beta^h \frac{u'_i(C_{i,t+h})}{u'_i(C_{i,t})} 1_{\{t+h \leq T\}}. \quad (3)$$

Hence, if agents expect AI to lead to an acceleration in the rate of aggregate consumption growth this will, ceteris paribus, lead to a drop in the SDF and additional discounting of future payoffs. Since u is concave, however, even news which increases expected future consumption $E_t[C_{t+h}]$ could lead to an increase in the mean of the SDF and thus a *decrease* in discounting if it implies a sufficiently large increase in uncertainty.

Bond Pricing Implications While the analysis above applies to general payoff streams Y , our empirical analysis will focus on bond prices. To study the implications of AI beliefs for bond prices, let 1_{t+h} denote a risk-free, h -period-ahead zero-coupon bond (i.e. the risk-free bond which pays one unit of consumption h periods in the future and nothing at any other time). By Equation (1) this bond's time- t price is given by

$$V_t(1_{t+h}) = \mathbb{E}_t[M_{t,t+h}] = \mathbb{E}_t \left[\beta^h \frac{u'(C_{t+h})}{u'(C_t)} \mathbf{1}_{\{t+h \leq T\}} \right],$$

and so is simply the expected h -period-ahead SDF. Since it is more common to work with bond yields than with prices, note that the period- t yield on the risk-free bond 1_{t+h} can be written as $y_{t,t+h} \equiv V_t(1_{t+h})^{-\frac{1}{h}} = \mathbb{E}_t[M_{t,t+h}]^{-\frac{1}{h}}$, which can be further re-written as

$$y_{t,t+h} = \frac{1}{\beta \mathbb{P}(t+h \leq T)^{\frac{1}{h}} \mathbb{E}_t \left[\frac{u'(C_{t+h})}{u'(C_t)} \mid t+h \leq T \right]^{\frac{1}{h}}}.$$

Thus, zero-coupon bond yields are decreasing in the discount factor β , increasing in the probability that T arrives before the bond pays off $\mathbb{P}(t+h > T)$, and decreasing in expected marginal utility in period $t+h$ conditional on T not yet having arrived, $\mathbb{E}_t \left[\frac{u'(C_{t+h})}{u'(C_t)} \mid t+h \leq T \right]$. Since u is concave, yields are thus increasing in consumption growth. Thus, as noted by Chow et al. (2024) both an increase in anticipated consumption growth and a closer expected arrival for T lead to higher risk-free yields.³

While the Treasuries which are the focus of our analysis below are multi-period rather than zero-coupon bonds, the comparative statics are much the same. In particular, if we consider a h -period risk-free bond with coupon c and face value d , this corresponds to

³We note, however, that beliefs about consumption growth and about T have distinct implications for the prices of *risky* assets. In particular, if we consider the ratio of risky and risk-free asset prices for a given future period, $V_t(Y_{t+h})/V_t(1_{t+h}) = \mathbb{E}_t \left[Y_{t+h} \frac{u'(C_{t+h})}{u'(C_t)} \mid t+h \leq T \right]$, this ratio depends only on behavior conditional on T not yet having arrived. This fact may be useful for distinguishing changes in beliefs about consumption from changes in beliefs about T .

payoff stream $B = \{c1_{t+1}, c1_{t+2}, \dots, c1_{t+h-1}, (c+d)1_{t+h}\}$ and so has price

$$V_t(B) = d\mathbb{E}_t[M_{t,t+h}] + c \sum_{s=1}^h \mathbb{E}_t[M_{t,t+s}] = \frac{d}{y_{t,t+h}^h} + c \sum_{s=1}^h \frac{1}{y_{t,t+s}^s}.$$

Empirical Strategy The model above suggests an empirical strategy for learning about changes in AI beliefs from asset prices: if we have a date t at which we believe information arrived about the future course of AI, then changes in long-dated asset prices around this date should incorporate the impact of the new information about AI.

To fix ideas, again consider the price for an asset that pays Y_{t+h} units in period $t+h$. If we think new information about AI arrived at t , we may compare prices at t_- and t_+ for $t_- < t < t_+ \ll h$, and use the fact that $V_{t_-}(Y_{t+h}) = \mathbb{E}_{t_-}[M_{t_-,t_+} V_{t_+}(Y_{t+h})]$ to write

$$V_{t_+}(Y_{t+h}) - V_{t_-}(Y_{t+h}) = V_{t_+}(Y_{t+h}) - \mathbb{E}_{t_-}[V_{t_+}(Y_{t+h})] - \mathbb{E}_{t_-}[(M_{t_-,t_+} - 1)V_{t_+}(Y_{t+h})].$$

If the time difference $t_+ - t_-$ is reasonably small we expect the final term to be negligible.⁴ Hence, by (1) and the law of iterated expectations we can approximate

$$V_{t_+}(Y_{t+h}) - V_{t_-}(Y_{t+h}) \approx \mathbb{E}_{t_+}[M_{t_+,t+h} Y_{t+h}] - \mathbb{E}_{t_-}[M_{t_+,t+h} Y_{t+h}]. \quad (4)$$

Thus the change in prices between t_- and t_+ gives us, approximately, the difference in conditional expectations for the discounted payoff $M_{t_+,t+h} Y_{t+h}$ at information sets \mathcal{F}_{t_-} and \mathcal{F}_{t_+} . In particular, if we consider the risk-free asset $Y_{t+h} = 1_{t+h}$, changes in prices reveal the change in the conditional mean of the SDF $M_{t_+,t+h}$.

For a given pair t_- and t_+ the difference $V_{t_+}(Y_{t+h}) - V_{t_-}(Y_{t+h})$ reflects all information that arrives between those dates, not just information about AI. Hence, in our empirical analysis we will aggregate across a series of AI news dates. So long as there is not other price-relevant information which systematically arrived at the same time as AI news, comparing behavior at AI dates to that at other dates will isolate the effect of AI news, though it will be important to account for the possibility of other news when assessing statistical uncertainty. As already noted, we will also use data on multi-period bonds rather than zeros. Since our primary focus will be on long-maturity bonds, however, most bond payoffs will be in the future and the intuition provided above for zeros will again translate to the bonds we study.

⁴By Cauchy-Schwarz, $\mathbb{E}_{t_-}[(M_{t_-,t_+} - 1)V_{t_+}(Y_{t+h})] \leq \sqrt{\mathbb{E}_{t_-}[(M_{t_-,t_+} - 1)^2]} \sqrt{\mathbb{E}_{t_-}[V_{t_+}(Y_{t+h})^2]}$.

3 Data and Methods

As the theory above suggests, if market participants think that AI may have large growth effects then new information about the trajectory of AI should impact long-term asset prices, including for assets that are not directly related to AI such as long-term risk-free bonds. We examine this prediction empirically, describing the data and methods we employ in this section and our empirical results in the next.

3.1 AI News Events

To look for asset prices changes around the arrival of AI news, we need to know a set of dates at which AI information arrived. While there are a variety of reasonable approaches one might take to this problem, we focus on release dates for new generative AI models from five major AI laboratories: OpenAI, Google DeepMind, Anthropic, xAI, and DeepSeek.⁵ For each lab, we focus on major updates to the lab’s flagship model series (e.g. ChatGPT in the case of OpenAI), and use the release date from the lab’s website.⁶ We limit attention to releases in calendar years 2023 and 2024, a period that (i) follows the November 2022 release of ChatGPT, which saw a significant increase in public attention to AI capabilities, and (ii) precedes the tariff announcements and other US macroeconomic policy changes that began in 2025. For OpenAI we include the “reasoning model” o1, since other labs included such models as part of their flagship series rather than numbering them separately (e.g. Gemini 2.5 from Google and Claude 3.7 from Anthropic, both released after our main analysis window). Table 1 collects the resulting release dates.

We use AI model releases as our event dates in order to capture new, forward-looking information about AI capabilities, rather than other aspects of technology or financial performance of firms. Put differently, our hypothesis is that major model releases provide information not only about the current state of AI capabilities but also about the rate of technological progress, potentially causing market participants to update their beliefs about future AI development. These events are also less directly linked to financial outcomes than some other plausible event dates, such as earnings announcements. At the same time, it is clear that information about AI system capabilities arrives outside of new model releases for these particular AI labs. There are many other AI researchers and firms, and even the

⁵These are the laboratories appearing in the top 10 style-adjusted rankings on the Chatbot Arena leaderboard as of June 29, 2025 (Chiang et al., 2024).

⁶For DeepSeek V2, we were unable to find an announcement on the lab’s website, and so instead use an announcement date from DeepSeek’s X account.

Table 1: AI Model Release Dates

Date	AI Laboratory	Model
<i>2023 Releases</i>		
02/06/2023	Google	Bard
03/14/2023	OpenAI	ChatGPT 4
03/14/2023	Anthropic	Claude 1
07/11/2023	Anthropic	Claude 2
11/03/2023	xAI	Grok 1
11/21/2023	Anthropic	Claude 2.1
12/06/2023	Google	Gemini Pro 1.0
<i>2024 Releases</i>		
02/15/2024	Google	Gemini Pro 1.5
03/04/2024	Anthropic	Claude 3
03/28/2024	xAI	Grok 1.5
05/06/2024	DeepSeek	DeepSeek V2
05/13/2024	OpenAI	ChatGPT 4-o
06/20/2024	Anthropic	Claude 3.5 Sonnet
08/13/2024	xAI	Grok 2
09/05/2024	DeepSeek	DeepSeek 2.5
12/05/2024	OpenAI	o1
12/11/2024	Google	Gemini 2.0
12/26/2024	DeepSeek	DeepSeek V3

Notes: This table presents the major AI model releases used in our event study analysis.

firms we study make numerous announcements and incremental model releases outside the set of major releases we consider. So long as some information is arriving around the dates we study, such alternative information sources do not pose a threat to the validity of our estimates, though as we discuss in Section 6 below it may matter for interpretation.

More directly relevant for us, for at least some model releases we know that certain experts were given early access to the model prior to the official release.⁷ To partially capture such information “leakage” our empirical specifications will include a window of dates prior to the model release (15 sample days, or approximately 3 weeks, for our preferred specifications). While this extended window is still unlikely to capture all information leakage, uncaptured leakage should reduce the amount of information arriving in our event windows. We expect this will bias us against finding yield responses.

⁷See for instance Mollick (2024).

3.2 Financial Market Data

Motivated by the theory in Section 2, to look for effects of AI information on long-run consumption expectations we examine the behavior of bond yields of different maturities around major model release dates. We consider three bond series.

1. **Nominal Treasury Yields:** We use constant-maturity Treasury yields from the Federal Reserve Economic Data (FRED) database for maturities of 1, 5, 10, 20, and 30 years (Board of Governors of the Federal Reserve System, US, 2025b).
2. **TIPS Yields:** We use constant-maturity Treasury Inflation Protected Securities (TIPS) yields from FRED for maturities of 5, 10, 20, and 30 years (Board of Governors of the Federal Reserve System, US, 2025c).
3. **Corporate Bond Indices:** We use ICE BofA corporate bond effective yield indices broken out by maturity (1-3 year, 3-5 year, 5-7 year, 7-10 year, 10-15 year, and 15+ year indices – Ice Data Indices, LLC 2025a).

All yield data are measured in percentage points and recorded at daily frequency. Since the available dates vary slightly, when analyzing each series we use all dates for that series in the analysis window.

3.3 Event Study Methodology

We use an event study approach to look for changes in yields around our event dates. For each AI event date $t \in \mathcal{T}$ and each yield series, we calculate the change in yields relative to a pre-event baseline, defined as b days before the event. Thus, the change from the baseline date to relative date s is:

$$\Delta y_{t,s} = y_{t+s} - y_{t-b}. \quad (5)$$

This gives us a yield change for each event date $t \in \mathcal{T}$. We next aggregate these changes across event dates to obtain a single summary statistic, considering both the median change

$$\text{MedianChange}_s = \text{Median}(\Delta y_{t,s}, t \in \mathcal{T})$$

and the median absolute change

$$\text{MedianAbsChange}_s = \text{Median}(|\Delta y_{t,s}|, t \in \mathcal{T})$$

The median measures whether there were systematic patterns in the direction of yield changes around our event dates, while the median absolute change measures whether there were systematic patterns in the magnitude of yield changes. We focus on medians, rather than means, because medians are more robust to outliers, which we view as especially important given our small sample size. Appendix B.5 provides results for mean and mean absolute changes, which prove to be qualitatively similar to our main results.

3.4 Permutation Inference

To gauge whether markets are responding to AI model releases, we need a way to judge whether the yield movements we observe around model releases are larger than one would expect due to chance. Given our very limited sample size, it is important to use a method that is valid in small samples. To this end, we assess statistical significance via permutation inference, under the assumption that our AI release dates are as good as randomly assigned and, in particular, can be treated as a uniform random draw from the dates in our analysis window.

Our procedure works as follows:

1. We define the set of potential “placebo” event dates as all days in our sample (subject to the full event window from $t_- = t - b$ to $t_+ = t + s$ being within the sample).
2. For each $m \in \{1, \dots, 5000\}$ we randomly sample (without replacement) K placebo dates from this set, where K equals the number of actual AI events in our sample (again restricted to events where the full event window is within the sample), and compute our test statistics using these placebo dates.
3. We compare the test statistics computed using the actual model release dates to the empirical distribution across placebo samples. If markets did not react to AI events, the event dates were selected as good as randomly, and yields were continuously distributed, then the probability that our observed test statistics would exceed the p -th percentile of the placebo distribution would equal p up to simulation error. In reality our yield data are only measured up to the level of basis points, so in cases of ties we round away from statistical significance.

Table 2: Two-sided p-values based on constant-maturity US Treasury yields

Maturity	Median Change		Median Absolute Change	
	± 5 days	± 15 days	± 5 days	± 15 days
1 Year	0.369	0.231	0.729	0.682
5 Year	0.189	0.231	1.000	0.798
10 Year	0.120	0.150	0.994	0.958
20 Year	0.097*	0.054*	0.867	0.563
30 Year	0.064*	0.038**	0.806	0.764

Notes: The “Median Change” columns consider the median change in yields across event dates, while the “Median Absolute Change” columns consider median absolute changes. For each statistic, we compare yields 5 or 15 days before and after each model release (in ± 5 the ± 15 columns, respectively). P-values are computed based on drawing placebo event dates 5000 times (uniformly at random from days in the sample with sufficient window on either side) and comparing resulting placebo distributions to observed changes around AI model releases. ** (*) denotes statistical significance at the 5% (10%) level.

This approach gives a test for the “sharp” null hypothesis of no impact on yields from any release, which is valid in finite samples under the auxiliary assumption that the AI model release dates can be treated as a random draw. While this is a strong assumption, it may be justified if model releases are driven by technical development timelines rather than financial market conditions. Examining the release dates in Table 1 we do not observe very strong calendar patterns, though e.g. Fridays appear somewhat underrepresented (with only 1 of the 17 unique dates in the sample), and there are more dates in 2024 than 2023 (with 11 of the 17 unique dates). If one wanted to replace our assumption that release dates are drawn uniformly at random with some other specific assumption about their distribution, our approach generalizes directly. To explore sensitivity to our assumptions, we discuss several robustness checks following our empirical results.

4 Empirical Results

We next report our empirical results. We begin by examining whether there are statistically significant changes in yields around our event dates, evaluating statistical significance relative to the placebo distribution as described in the previous section.

Recall that p-values measure the probability of observing a more extreme outcome were the null hypothesis true. Hence, small p-values correspond to outcomes which are

unlikely to arise under the null (in our case, if AI model releases have no effect on yields, and release dates are as good as random). Consequently, a 10% test of the null rejects when the p-value is less than 0.1, and a 5% test rejects when the p-value is less than 0.05. Tables 2-4 report two-sided p-values for the fixed-income yield series we consider (US Treasuries, TIPS, and corporate bond indices), reporting results for both median and median absolute changes, and comparing yields either five or fifteen days before and after each model release (that is, setting $b=s=5$ or $b=s=15$ in the notation of Equation 5).

Table 3: Two-sided p-values based on constant-maturity TIPS yields

Maturity	Median Change		Median Absolute Change	
	± 5 days	± 15 days	± 5 days	± 15 days
5 Year	0.341	0.128	0.549	0.576
10 Year	0.182	0.107	0.262	0.601
20 Year	0.114	0.096*	0.350	0.756
30 Year	0.112	0.038**	0.257	0.783

Notes: The “Median Change” columns consider the median change in yields across event dates, while the “Median Absolute Change” columns consider median absolute changes. For each statistic, we compare yields 5 or 15 days before and after each model release (in ± 5 the ± 15 columns, respectively). P-values are computed based on drawing placebo event dates 5000 times (uniformly at random from days in the sample with sufficient window on either side) and comparing resulting placebo distributions to observed changes around AI model releases. ** (*) denotes statistical significance at the 5% (10%) level.

The results in Tables 2-4 paint a consistent picture. First considering median changes in bond yields, we see evidence of changes in yields for longer maturity bonds for the ± 5 day specifications, though the p-values sometimes fall short of significance at conventional levels. For the ± 15 day specifications, we see statistically significant changes in yields for longer-maturity bonds. This holds true whether we consider Treasuries, TIPS, or corporate bonds. By contrast, when we consider median absolute changes we do not find statistically significant effects at conventional significance levels for any of the maturities studied. These patterns again hold across Treasuries, TIPS, and corporate bonds.

This pattern is different than we, at least, anticipated before analyzing the data: if market participants took seriously the possibility of transformative AI, and learned more than usual about AI’s future trajectory around model release dates, we would expect larger-than-average yield changes around model release dates (and hence, potentially, statistical significance for median absolute changes) but not necessarily a consistent direction of change

Table 4: Two-sided p-values based on ICE corporate bond index yields

Maturity	Median Change		Median Absolute Change	
	± 5 days	± 15 days	± 5 days	± 15 days
1-3 Year	0.531	0.086*	0.722	0.827
3-5 Year	0.029**	0.037**	0.831	0.905
5-7 Year	0.036**	0.036**	0.654	0.963
7-10 Year	0.055*	0.040**	0.443	0.853
10-15 Year	0.049**	0.046**	0.472	1.000
15+ Year	0.100*	0.051*	0.864	0.993

Notes: The “Median Change” columns consider the mean change in yields across event dates, while the “Median Absolute Change” columns consider median absolute changes. For each statistic, we compare yields 5 or 15 days before and after each model release (in ± 5 the ± 15 columns, respectively). P-values are computed based on drawing placebo event dates 5000 times (uniformly at random from days in the sample with sufficient window on either side) and comparing resulting placebo distributions to observed changes around AI model releases. ** (*) denotes statistical significance at the 5% (10%) level.

(and hence, potentially, no statistical significance for median changes). Our results show the opposite: there do not appear to be yield changes of statistically different magnitude around AI model release dates (since we do not see statistical significance for median absolute changes). However, there are statistically significant patterns in the direction of yield changes, especially at longer maturities, as revealed by our results on median changes.

Event Study Plots To further explore what is happening around our AI events, Figures 2-4 plot, for each yield series and each horizon $s \in \{-14, \dots, 15\}$, the median and median absolute change in yields (relative to $b = 15$ days before the event) across the observed AI model releases.⁸ For comparison, at each horizon we also plot the mean of the placebo distribution and bands which contain, 90%, 95%, and 99% of the placebo draws pointwise at each horizon (with equal mass assigned to the two tails). These bands are another way to express our placebo tests. For instance, our placebo test rejects the null of no effect at the 10% level at a given horizon if and only if the median change at that horizon lies outside the 90% placebo band.

Examining these plots, we see that for both median and median absolute changes there

⁸To hold the set of events constant across different horizons, we limit attention to those model releases where the full window $[t-14, t+15]$ falls in calendar years 2023-4. This corresponds to the first 15 releases in Table 1, dropping Gemini 2.0 and DeepSeek V3.




Figure 2: Median and median absolute change in yields (relative to fifteen days before event) for constant-maturity US Treasury Bonds. Median taken across AI release events in the 2023 and 2024 calendar years. Placebo distribution recomputes statistics on dates drawn uniformly at random from sample period.




Figure 3: Median and median absolute change in yields (relative to fifteen days before event) for constant-maturity inflation-protected US Treasury Bonds (TIPS). Median taken across AI release events in the 2023 and 2024 calendar years. Placebo distribution recomputes statistics on dates drawn uniformly at random from sample period.




Figure 4: Median and median absolute change in yields (relative to fifteen days before event) for corporate bond indices. Median taken across AI release events in the 2023 and 2024 calendar years. Placebo distribution recomputes statistics on dates drawn uniformly at random from sample period.

is limited (and largely statistically insignificant) departure from the placebo distribution between $t-15$ and $t-5$. Bond yields, especially for longer-maturity bonds, show declines starting between $t-5$ and $t-2$. These declines continue through at least $t=0$, and lower yields persist through $t+15$. The apparent anticipatory effects (i.e. effects before the model release date t) are consistent with the fact, discussed above, that some information about new models may become available to market participants prior to the official model release.

The overall fall in yields around model releases is quantitatively large, exceeding 10 basis points by the end of the window for most series. Moreover, consistent with our previous findings these changes are statistically significant relative to the placebo distribution at conventional significance levels. Thus, we find economically and statistically significant declines in long-maturity bond yields around AI model releases, where these declines persist for at least three weeks after the release date.

Corporate Bond Spreads Figures 2 and 4 show a significant decline in both Treasury and corporate yields around AI model release dates, especially at the long end of the yield curve. These observations raise an immediate question: is there any change in corporate yields above and beyond the change in Treasury yields? Put differently, is the impact on the corporate yield curve fully explained by the change in Treasury yields, or does AI news have an additional impact on corporate bond yields?

To answer this question, Figure 5 plots the event study for the ICE BofA Option-Adjusted Spread index, where spreads are measured relative to US Treasuries (Ice Data Indices, LLC, 2025b). Comparing the observed changes in spreads to placebo bands we find no statistically significant changes in spreads.⁹

Exchange Rates Given our findings on bond yields, one might wonder whether AI model releases are leading to international capital flows. To provide some evidence on this point, in Appendix A we plot the event study for a broad trade-weighted US dollar exchange rate index around our model release dates (Board of Governors of the Federal Reserve System, US, 2025d). We find that AI model releases are associated with a statistically significant depreciation of the dollar, which starts a few days before the model release and persists through 15 days after. These declines are more gradual than the bond yield changes we find above, but appear consistent with e.g. a depreciation of the dollar following a drop in interest rates.

⁹We also find no significant effects on spreads when we look at corporate bond indices broken out by credit rating, though for the sake of brevity we do not report those results.




Figure 5: Median change in option-adjusted spreads (relative to fifteen days before event) for corporate bond indices. Median taken across AI release events in the 2023 and 2024 calendar years. Placebo distribution recomputes statistics on dates drawn uniformly at random from sample period.

4.1 Robustness Checks

We conduct a variety of analyses to explore the robustness of our result. Here we briefly discuss these robustness checks, presenting all results, along with additional details, in Appendix B. For brevity, in these robustness checks we primarily focus on results for US Treasuries.

Robustness to Dropping Events Since we examine yield changes around a relatively small number of model releases, one might worry that our findings could be driven by one or a few extreme events. For instance, the March 14, 2023 model releases in our data occurred soon after the March 10 collapse of Silicon Valley Bank. Our focus on medians rather than means is intended to mitigate this and similar concerns involving a small number of dates. To verify robustness of our results, Appendix B.1 reports versions of our results when we drop all subsets

of one, two, and three dates from our AI model release date series.¹⁰ We find that our results are quantitatively similar, and retain statistical significance at many horizons, when dropping any one date from our sample. Even when dropping two or three dates our results remain directionally similar, and are robustly significant (at the 10% level) at certain horizons.

Alternative “Placebo” Dates Our hypothesis tests and p-values are based on the assumption that AI model release dates are as good as random and, consequently, that systematic moves in bond yields around AI model releases may be attributed to beliefs about AI. Our inference results would thus be invalid if the timing of AI model releases were systematically related to yield movements for other reasons, for instance because AI labs attempt to time their releases around market movements directly, or because they time model releases around other, non-AI events which systematically move markets. While our results above show that our findings are robust to dropping any small set of “suspect” model releases, they do not address the possibility of more pervasive timing correlation.

For any alternative date series, an extreme form of timing correlation would be for AI model releases to be drawn solely from that series. If the subset of dates selected were as good as random from within that series, we could repeat our placebo calculation to derive thresholds for statistical significance. Motivated by this observations, Appendix B.2 reports versions of the median change plots in Figure 2 which use one of (i) FOMC meetings (ii) major tech-firm annual events (iii) major tech firm earnings releases, (iv) CPI release dates, (v) jobs report release dates (vi) retail sales release dates, and (vii) Treasury auction dates for 10, 20, and 30 year bonds as the source of our placebo dates, though in fact none of these series nests our AI model release series. Our findings remain statistically significant relative to these alternative “placebo” distributions.

Controlling for Other News As a further robustness check, we directly control for proxies for certain non-AI news that arrived during our analysis period. Specifically, we consider three series intended to capture other information that might have impacted bond yields (i) the Citigroup US Economic Surprise Index (Citigroup Global Markets 2025, which summarizes the deviation of economic data releases from forecasts) (ii) the Cboe’s VIX volatility index (Cboe 2025, which is an option-implied measure of stock market volatility), and (iii) the Federal Reserve Bank of San Francisco Daily News Sentiment Index (Shapiro

¹⁰We drop event dates, rather than model releases, so dropping the two model releases on March 14, 2023 “counts” as dropping a single date.

et al. 2022; Federal Reserve Bank of San Francisco 2023, which summarizes the economic sentiment of news articles from a variety of sources). In each case, and for each of the US Treasury series we consider, we residualize daily changes in yields against the current level and 15 daily lags of the “control” series, then repeat our analysis with the re-cumulated series (now testing the null of no effect on the residualized yield series). Appendix B.3 shows that our results are directionally similar, and statistically significant at some horizons, whether controlling for any of the individual series or all three at once, though the level of significance varies across specifications.

Results for Alternative Analysis Samples Appendix B.4 reports versions of our Treasury results for alternative analysis samples, first plotting results for calendar years 2023 and 2024 separately, and then plotting results for an extended sample period running from October 2022 through May 2025. We find directionally similar results in all cases, though the results for 2023 are quantitatively larger than our main results, while those for 2024 are only marginally statistically significant at intermediate horizons, and lose statistical significance at longer horizons. The results for the extended sample period are similar to our main results.

Means vs. Medians Finally, Appendix B.5 reports versions of our main results instead considering mean and mean absolute changes. Our findings there are similar to those reported above.

5 Interpretation

Our empirical analysis shows that major AI model releases were accompanied by reductions in long-term bond yields. As discussed in Section 2, viewed through the lens of the complete-market, representative agent model, falling yields on the risk-free asset imply that the expected future marginal utility of consumption is rising, because expected future consumption is falling, uncertainty is increasing, or the date T after which asset holdings are irrelevant is believed to be shifting further into the future (or is less likely to arrive at all).

One natural question, in light of our findings, is how much investors must have updated their beliefs about growth in order to rationalize observed changes in yields. Providing a quantitative answer to this question requires imposing additional assumptions beyond those in Section 2. Since this interpretive exercise nevertheless appears worthwhile, in Section 5.1 we consider a more restrictive version of our model which we use to quantitatively

interpret our findings.

A second natural question is how to interpret investors' updated beliefs. In particular, to the extent investors lowered their consumption growth expectations around the model releases in our sample, does this reflect that they were positively surprised by the rate of AI progress and thought AI would be bad for consumption growth? Or did they think AI would be good for consumption growth but that the rate of AI progress was disappointing? Section 5.2 provides suggestive evidence on this point, using data from an online prediction platform to show that platform participants were, on median across the model releases in our sample, positively surprised by the rate of AI progress on at least some dimensions.

5.1 A Simplified Model

As discussed in Section 2, the assumption of complete markets implies the existence of a representative agent, so in this section we focus on that agent's consumption and utility. We assume that the representative agent has CRRA flow utility from consumption, $u(C_t) = \frac{C_t^{1-\gamma}}{1-\gamma}$, similar to Jones (2024).¹¹ Under this assumption, the SDF simplifies to

$$M_{t,t+h} = \beta^h \left(\frac{C_{t+h}}{C_t} \right)^{-\gamma} 1_{\{t+h \leq T\}}.$$

Unfortunately this restriction does not, on its own, suffice to let us interpret our empirical findings, since the consumption process C_t may have quite rich dynamics, reflecting many factors other than AI. To isolate the impact of AI beliefs, we thus impose further assumptions which restrict the evolution of C_t , and beliefs about C_t , over time. As a starting point, we assume there exists a horizon $k \geq 0$ such that at each date t in our sample the representative agent thinks that for all horizons $h \geq k \geq 0$ periods in the future, aggregate consumption evolves according to

$$C_{t+h+1} = (1+g)X_{t+h+1}C_{t+h}$$

where g captures the consumption growth impact of AI and $\{X_s\}_{s=t+k+1}^{\bar{T}}$ is a stochastic process capturing the non-AI determinants of consumption growth. We make this assumption starting k periods in the future, rather than immediately, to allow the possibility of richer dynamics in short-term consumption, e.g. because the growth impacts of AI could

¹¹Jones (2024) includes an additional location term in the utility, but this term will be irrelevant for our purposes so we drop it.

take some time to “kick in.” Together with CRRA utility, this implies that for $h \geq k$ the h -period ahead SDF is

$$M_{t,t+h} = \left(\frac{C_{t+k}}{C_t} \right)^{-\gamma} \beta^h (1+g)^{-(h-k)\gamma} \left(\prod_{s=k+1}^h X_{t+s} \right)^{-\gamma} 1\{t+h \leq T\}. \quad (6)$$

We further assume that conditional on information available at t and the event $t+k \leq T$, (i) $\left(\{X_s\}_{s=t+k+1}^{t+h}, \left(\frac{C_{t+k}}{C_t} \right)^{-\gamma} \right)$, T , and g are believed to be mutually independent (ii) T is thought to arrive with probability δ_t in each period following $t+k$,

$$\mathbb{P}_t(t+h \leq T | t+k \leq T) = \prod_{s=k+1}^h \mathbb{P}_t(t+s \leq T | t+s-1 \leq T) = (1-\delta_t)^{h-k},$$

and (iii) $1+g$ is believed to be log-normally distributed, $\log(1+g) | \mathcal{F}_t, t+k \leq T \sim N(\mu_t, \sigma_t^2)$.¹²

These assumptions, taken together, imply tractable expressions for log forward rates which may in turn be used to interpret our empirical results. Consider the the period t forward yield from $t+k$ to $t+h$, i.e. the per-period yield earned by, in period t , selling a period $t+k$ zero-coupon risk-free bond while buying a period $t+h$ zero,

$$f_{t+k,t+h} = \left(\frac{y_{t,t+h}^h}{y_{t,t+k}^k} \right)^{\frac{1}{h-k}}.$$

Appendix C.1 shows that under our assumptions (i)-(iii) above, the log forward yield is

$$\begin{aligned} \log(f_{t+k,t+h}) &= \frac{1}{h-k} \log \left(\frac{\mathbb{E}_t[M_{t,t+k} | t+k \leq T]}{\mathbb{E}_t[M_{t,t+h} | t+k \leq T]} \right) = \frac{h}{h-k} \log(y_{t,t+h}) - \frac{k}{h-k} \log(y_{t,t+k}) = \\ & -\log(\beta) - \log(1-\delta_t) + \gamma\mu_t - \frac{\gamma^2}{2}(h-k)\sigma_t^2 - \frac{1}{h-k} \log \left(\mathbb{E}_t \left[\frac{\left(\frac{C_{t+k}}{C_t} \right)^{-\gamma}}{\left[\left(\frac{C_{t+k}}{C_t} \right)^{-\gamma} \right]} \left(\prod_{s=k+1}^h X_{t+s} \right)^{-\gamma} \right] \right). \end{aligned}$$

¹²These assumptions are restrictive, and appear unlikely to hold exactly. For instance, one might expect that more effective AI (i.e. AI yielding a higher g) would be associated with a closer arrival date for T . Similarly, if the growth effects of AI may “kick in” strictly before period $t+k$ then a higher g should lead to a higher C_{t+k} . Nevertheless, additional assumptions are needed to quantitatively interpret our results, and those above are the least objectionable assumptions we have thus far found that suffice to yield tractability.

Consequently, if we difference the log forward yields at two dates $t_- < t < t_+$ we have

$$\begin{aligned} & \log(f_{t_++k,t_++h}) - \log(f_{t_-+k,t_-+h}) = \\ & -\log\left(\frac{1-\delta_{t_+}}{1-\delta_{t_-}}\right) + \gamma(\mu_{t_+} - \mu_{t_-}) - \frac{\gamma^2}{2}(h-k)(\sigma_{t_+}^2 - \sigma_{t_-}^2) - \eta_{t_-,t_+,k,h} \end{aligned}$$

where

$$\eta_{t_-,t_+,k,h} = \log\left(\frac{\mathbb{E}_{t_+}\left[\left(\frac{C_{t_++k}}{C_{t_+}}\right)^{-\gamma}\left(\prod_{s=k+1}^h X_{t_++s}\right)^{-\gamma}\right]}{\mathbb{E}_{t_-}\left[\left(\frac{C_{t_-+k}}{C_{t_-}}\right)^{-\gamma}\left(\prod_{s=k+1}^h X_{t_-+s}\right)^{-\gamma}\right]}\right) \cdot \frac{\mathbb{E}_t\left[\left(\frac{C_{t_-+k}}{C_{t_+}}\right)^{-\gamma}\right]}{\mathbb{E}_t\left[\left(\frac{C_{t_++k}}{C_{t_+}}\right)^{-\gamma}\right]}.$$

To connect this expression to our empirical results, let us again consider our set of event dates $t \in \mathcal{T}$, and for each t consider $t_+ = t + s$ and $t_- = t - b$.¹³ Let \mathcal{A} denote the set of all dates t such that t_+ and t_- are both in the sample. We assume that for all $h \geq k$ the residuals $\eta_{t_-,t_+,k,h}$ have approximately the same mean across our event dates \mathcal{T} as across \mathcal{A} ,

$$\frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \eta_{t_-,t_+,k,h} \approx \frac{1}{|\mathcal{A}|} \sum_{t \in \mathcal{A}} \eta_{t_-,t_+,k,h}. \quad (7)$$

For instance, if we assumed that $\eta_{t_-,t_+,k,h}$ were stationary across time conditional on our event dates \mathcal{T} and regularity conditions held, this would follow from the law of large numbers as $|\mathcal{T}| \rightarrow \infty$.¹⁴ Motivated by this assumption, we consider the difference in differences of log forward rates $\log(f_{t+k,t+h})$ across times $t \in \mathcal{T}$ and $t \in \mathcal{A}$:

$$\begin{aligned} \text{DID}(\log(f_{t+k,t+h}); \mathcal{T}, \mathcal{A}) & \equiv \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \log\left(\frac{f_{t_++k,t_++h}}{f_{t_-+k,t_-+h}}\right) - \frac{1}{|\mathcal{A}|} \sum_{t \in \mathcal{A}} \log\left(\frac{f_{t_++k,t_++h}}{f_{t_-+k,t_-+h}}\right) \approx \quad (8) \\ & \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \left(-\log\left(\frac{1-\delta_{t_+}}{1-\delta_{t_-}}\right) + \gamma(\mu_{t_+} - \mu_{t_-}) - \frac{\gamma^2}{2}(h-k)(\sigma_{t_+}^2 - \sigma_{t_-}^2) \right) - \\ & \frac{1}{|\mathcal{A}|} \sum_{t \in \mathcal{A}} \left(-\log\left(\frac{1-\delta_{t_+}}{1-\delta_{t_-}}\right) + \gamma(\mu_{t_+} - \mu_{t_-}) - \frac{\gamma^2}{2}(h-k)(\sigma_{t_+}^2 - \sigma_{t_-}^2) \right) \end{aligned}$$

¹³Thus, t_+ and t_- are implicitly functions of t , though we suppress this dependence for readability.

¹⁴We work with means across event dates, rather than medians as in Section 4, because means recover a simple aggregation of heterogeneous effects across events, while we are not aware of a similarly tractable expression for medians.

Thus, if we consider the slope of $\text{DID}(\log(f_{t+k,t+h}))$ with respect to the horizon h , this approximately recovers the difference in differences for the variance σ_t^2 , scaled by $-\frac{\gamma^2}{2}$.

$$-\frac{\gamma^2}{2}\text{DID}(\sigma_t^2; \mathcal{T}, \mathcal{A}) \equiv -\frac{\gamma^2}{2} \left(\frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} (\sigma_{t+}^2 - \sigma_{t-}^2) - \frac{1}{|\mathcal{A}|} \sum_{t \in \mathcal{A}} (\sigma_{t+}^2 - \sigma_{t-}^2) \right). \quad (9)$$

If we think that dates in $\mathcal{A} \setminus \mathcal{T}$ have little news relevant to the growth impacts of AI, we might expect the second term to be small relative to the first. However, our event dates are also included in the second term and we moreover do not want to rule out the possibility that AI-relevant news arrives at dates outside of \mathcal{T} . Hence, we focus on the difference-in-differences interpretation.

Similarly, the intercept of $\text{DID}(\log(f_{t+k,t+h}))$ as $h \downarrow k$ measures the difference in differences for expected log growth, scaled by γ , less the difference in differences in the log probability that T does not arrive in a given year,

$$\gamma \text{DID}(\mu_t; \mathcal{T}, \mathcal{A}) - \text{DID}(\log(1 - \delta_t); \mathcal{T}, \mathcal{A}) \quad (10)$$

for

$$\begin{aligned} \text{DID}(\mu_t; \mathcal{T}, \mathcal{A}) &\equiv \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} (\mu_{t+} - \mu_{t-}) - \frac{1}{|\mathcal{A}|} \sum_{t \in \mathcal{A}} (\mu_{t+} - \mu_{t-}) \\ \text{DID}(\log(1 - \delta_t); \mathcal{T}, \mathcal{A}) &\equiv \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \log\left(\frac{1 - \delta_{t+}}{1 - \delta_{t-}}\right) - \frac{1}{|\mathcal{A}|} \sum_{t \in \mathcal{A}} \log\left(\frac{1 - \delta_{t+}}{1 - \delta_{t-}}\right). \end{aligned}$$

Taking the Model to the Data Our simplified model predicts the behavior of yields on risk-free zero-coupon bonds, so to take these predictions to the data, we use daily Treasury yield curves from FRED (Board of Governors of the Federal Reserve System, US, 2025a), which are based on a three-factor term structure model due to Kim and Wright (2005). These data cover maturities up to 10 years.

To apply the above results, we must choose a horizon k beyond which to consider forward yields. To guide this choice, in Appendix C.2 we plot the difference in differences in one period-ahead log forward yields, $\text{DID}(\log(f_{t+h,t+h+1}); \mathcal{T}, \mathcal{A})$ for $h \in \{0, \dots, 9\}$. Equation (8) implies that for $h \geq k$ this curve should be approximately linear in h . This does not appear to hold exactly in our data, but for $h \geq 4$ it seems a good approximation. Motivated by this finding, for the remainder of our analysis we take $k = 4$.

After selecting $k = 4$, we regress $\text{DID}(\log(f_{t+k,t+h}); \mathcal{T}, \mathcal{A})$ for horizons $h \in \{5, \dots, 10\}$

on the difference $h - k$ relative to the initial horizon. This yields a slope of 0.0015 log points (corresponding to the 60.5th percentile of the placebo distribution) and an intercept of approximately -0.208 log points (corresponding to the 1.5th percentile of the placebo distribution).¹⁵ The finding that the slope of the yield curve is not substantially changing around model release dates is consistent with our finding in Section 4 above that the yield impacts of model releases are quite similar for the various bond maturities above 5 years. Thus, it appears that the changes we observe around event dates are driven by shifts in the level of the forward curve, rather than the slope.

To further interpret these results through the lens of the simplified model developed above, we separately consider the interpretation of the slope and intercept.

Interpreting the Slope First considering the slope (9) of the forward curve difference in differences (8) with respect to the maturity difference $h - k$, recall that the slope coefficient estimates the scaled average variance change. Thus, we estimate that the average change in the variance of $\log(1 + g_t)$ around model releases, less the variance change around the average date in the sample, is

$$\widehat{\text{DID}}(\sigma_t^2; \mathcal{T}, \mathcal{A}) = -\frac{3}{\gamma^2} \cdot 10^{-5},$$

for γ the CRRA coefficient of the representative agent. Hence, the simplified model considered in this section suggests that consumption growth uncertainty actually *fell* slightly on average around model release dates relative to the average day in our sample.

These estimates are small, and are not statistically different from zero according to our placebo distribution even using generous thresholds for statistical significance.¹⁶ This finding of little evidence for growth uncertainty changes around our event dates is consistent with our finding in Section 4 that there does not appear to be a clear trend in yield changes across 10, 20, and 30 year Treasuries. That said, given our limited sample size we do not have much power to detect small slope changes (the 5th and 95th percentiles of our placebo distribution correspond to slopes of approximately $\pm 8.9 \cdot 10^{-3}$, respectively).

Overall, our simplified model suggests that, if anything, consumption growth uncertainty

¹⁵A previous version of the paper mistakenly reported slope and intercept coefficients based on the extended 2022-5 sample considered in Appendix B.4, rather than our main 2023-4 sample.

¹⁶To interpret the magnitude of our estimated variance reduction, note that it is equivalent to, on the average event date, removing a noise component from $\log(1 + g_t)$ with standard deviation equal to $0.55/\gamma$ percentage points. While this is not a negligible uncertainty reduction for e.g. $\gamma \in [1, 5]$, it is delicate to interpret given its statistical insignificance.

may have slightly fallen around the model release dates we study, though our estimates are imprecise. Nevertheless, we have sufficient evidence to conclude that, through the lens of our simplified model, changes in consumption growth uncertainty do not explain the yield decreases we observe around AI event dates.

Interpreting the Intercept We next turn to the intercept (10) in the forward curve difference in differences (8). Recall that under our simplified model this term captures two forces: changes in the anticipated arrival rate δ_t of T (where a closer expected arrival for T increases yields) and changes in the mean μ_t of the log growth rate $\log(1+g_t)$ (where a higher value of μ_t again increases yields).

If our estimated intercept were due entirely to a change in beliefs about T , the model implies that the average model release in our sample led to a roughly 0.208 percentage point increase in $\log\left(\frac{1-\delta_{t+}}{1-\delta_{t-}}\right)$ relative to the average in the sample

$$\text{DID}(\widehat{\log(1-\delta_t)}; \mathcal{T}, \mathcal{A}) \approx 0.208\%.$$

If we assume δ_t is close to zero, it follows that $\text{DID}(\widehat{\delta_t}; \mathcal{T}, \mathcal{A}) \approx -0.208\%$, so we estimate that the average AI event in our sample is associated with a roughly 0.208 percentage point larger reduction in δ_t than the average date in the sample. Cumulated over the 15 model releases in our analysis sample, this corresponds to a 3.12 percentage point decrease in the annual arrival probability of T , which seems like a large effect.¹⁷

If observed changes in yields were instead due entirely to changes in consumption growth expectations, the model implies that the average model release in our sample led to an approximately $0.208/\gamma$ percentage point larger decrease in μ_t than the average date in the sample,

$$\text{DID}(\widehat{\mu_t}; \mathcal{T}, \mathcal{A}) = -\frac{0.208\%}{\gamma},$$

for γ the CRRA coefficient of the representative agent. If we assume that $\sigma_t^2 = \text{Var}_t(\log(1+g))$ is small for all t , and further assume that μ_t is close to zero, this implies that

$$\text{DID}(\widehat{\mathbb{E}_t[g]}; \mathcal{T}, \mathcal{A}) \approx -\frac{0.208\%}{\gamma},$$

¹⁷Direct adding-up of effects is complicated by the fact that the event windows for some of our model releases overlap. On the other hand, re-running our analysis on the extended sample considered in Appendix B.4 again produces large per-event effects, now over a larger set of events.




Figure 6: Values for $\text{DID}(\mu_t; \mathcal{T}, \mathcal{A}) \approx \text{DID}(\mathbb{E}_t[g]; \mathcal{T}, \mathcal{A})$ and $-\text{DID}(\log(1 - \delta_t); \mathcal{T}, \mathcal{A}) \approx \text{DID}(\delta_t; \mathcal{T}, \mathcal{A})$ compatible with an intercept value (10) equal to -0.208 percentage points under the simplified model and different levels of CRRA parameter γ .

so the average model release in our sample implies a roughly $0.208/\gamma$ percentage point reduction in expected consumption growth, relative to the average date in the sample. Thus, under $\gamma = 1$ (i.e. log utility) our results imply a 0.208 percentage point, or 20.8 basis point, drop in expected consumption growth (3.12 percentage points, cumulated), while under $\gamma = 2$ they imply a 0.104 percentage point drop (1.56 percentage points, cumulated), and under $\gamma = 5$ they imply a 0.041 percentage point drop (0.62 percentage points, cumulated). Even at the lower end, these again seem like substantial effects.

Of course, it could be that beliefs about both T and g update in response to AI model releases. To explore this broader range of possible interpretations, Figure 6 depicts the $\text{DID}(\mu_t; \mathcal{T}, \mathcal{A})$ and $-\text{DID}(\log(1 - \delta_t); \mathcal{T}, \mathcal{A}) \approx \text{DID}(\delta_t; \mathcal{T}, \mathcal{A})$ combinations compatible with an intercept (10) of -0.208 percentage points, for different levels of CRRA parameter γ . There is downward-sloping relationship between the implied effects on the arrival rate of T and g : the larger the decrease in the arrival rate of T , the more positive the growth effects which rationalize observed yield changes, and vice versa.

Overall, our simplified model implies that the changes in bond yields we observe around AI model release dates are primarily driven by some combination of decreases in growth expectations (i.e. μ_t) and decreases in the perceived arrival rate of T (i.e. δ_t) rather than changes in growth uncertainty (i.e. σ_t).

5.2 Suggestive Evidence on AI Belief Updating

To complement our results on bond yields, we next analyze AI-progress forecasts from the online prediction platform Metaculus. Metaculus is a forecasting platform where participants make probabilistic predictions about future events, with predictions aggregated to produce community forecasts. We focus on a Metaculus question regarding the arrival of “weakly general artificial intelligence,” or weak AGI, which asks users to predict the first date at which a unified AI system will be publicly known to satisfy a number of criteria (Metaculus, 2020b).¹⁸ A substantial number of participants contributed forecasts for this question, growing from over 600, at the start of our analysis window, to over 1500 by the end.

Metaculus provides a forecast distribution, based on weighted aggregation of individual participants’ forecast distributions, rather than simply a point forecast. We thus examine how the forecast distribution changes around our event dates, focusing on the 25th percentile, median, and 75th percentile of the forecast distribution, and taking the median change across event dates as for our main results.¹⁹ The results, shown in the first panel of Figure 7, show that the forecast distribution shifts down on median around model release dates in our sample, corresponding to an earlier arrival date for weak AGI and thus faster AI progress. The shifts in the 25th percentile and median are statistically significant at conventional significance levels, while that for the 75th percentile is marginally significant. Interestingly, as in our financial market results the downward shift in the 25th percentile of the forecast distribution occurs substantially before the model release, though the others occur later.

While we find these results interesting, they are sensitive to the precise question we consider. If we instead examine Metaculus’s question about the arrival of the first AGI system, which sets more demanding criteria than for weak AGI (Metaculus, 2020a), we do not see clear changes in the forecast distribution around model releases in our sample. If anything the forecast distribution increases, though these increases are largely statistically insignificant – see the second panel of Figure 7. One interpretation of these results could be that Metaculus participants thought the model releases we study were informative about

¹⁸Specifically, the weak AGI criteria involve: (1) scoring 90% or more on a robust version of the Winograd Schema Challenge, (2) scoring at the 75th percentile on the mathematics section of a circa-2015-2020 standard SAT exam, (3) passing a Turing test, and (4) learning to play the classic Atari game “Montezuma’s Revenge” based on less than 100 hours of real-time play. The question explicitly requires these capabilities be demonstrated by a unified system rather than separate specialized models cobbled together.

¹⁹Unlike our financial market data, Metaculus forecasts update every day, including on weekends, so our ± 15 day analysis window here corresponds to a shorter “real” time period than that for our other results. When computing the placebo distribution in this section, we restrict the placebo dates to be drawn from dates covered by our Treasuries series, since no model was released e.g. on a weekend or holiday.

the arrival of weak AGI, but that more fundamental progress is needed to attain AGI.²⁰

Overall, these results suggest that Metaculus participants updated positively about at least some aspects of AI progress around the AI model releases in our sample. While there is no guarantee that the beliefs of Metaculus participants resemble those of investors, to the extent the two are related these results suggest that the yield changes we observe around AI model release dates may not be driven by disappointing AI progress.

6 Discussion

We have found evidence of economically and statistically large declines in long-term bond yields around major AI model releases. Viewed through the lens of a simple asset pricing model, these results suggest that investors are updating their beliefs towards some combination of (i) lower expected consumption growth (ii) higher uncertainty about future consumption or (iii) a lower probability of extreme “doom” or “bliss” scenarios. We can roughly quantify the extent of belief updating under the additional assumptions laid out in Section 5. Since we find substantial shifts in the level, but not the slope, for the forward curve, the model implies that (i) and/or (iii) play a much more important role than (ii) in explaining our results.

These conclusions are subject to important caveats. Perhaps most important, it could be that the yield changes we observe around AI model releases do not reflect the causal effects of AI news and are instead driven by other factors. Even granting that the effects we estimate are causal, there are other possible interpretations. First, it may be that none of the bonds we consider is a reasonable proxy for a risk-free asset. Second, updates to investor beliefs around the model release dates we study could be non-representative of overall investor beliefs about AI, and third, the simple complete-market, representative agent model might imply a misleading interpretation of market responses. We discuss each possibility in turn.

On the first possibility, it is plausible that investors do not think US Treasuries are approximately risk-free. Treasuries are subject to inflation risk, and potentially to default risk given the large and growing budget deficits run by the US government. TIPS are designed to reduce inflation risk, and so partially mitigate this concern, but remain subject to default risk. If market participants think there is a non-trivial probability of a US default in the coming decades it could be that news about AI raises expected future tax

²⁰Consistent with this, the forecast median for AGI at the end of 2024 was August 2033, compared to March 2027 for weak AGI.




Figure 7: Median change in forecast distribution quantiles for weak AGI arrival date (top panel) and AGI arrival date (bottom panel). Median taken across AI release events in the 2023 and 2024 calendar years. Placebo distribution recomputes statistics on dates drawn uniformly at random from dates in US Treasury data series.

revenue, and thus lowers Treasury yields by lowering the embedded risk premium rather than by changing growth expectations.

Our data can provide some limited evidence on this possibility. One would expect that if the US government were to default this might increase the risk of many companies defaulting as well, so the mere fact that corporate bond yields also fall around AI model releases does not rule out this explanation. However, to the extent that not all highly profitable US companies would necessarily default if the US government did, we would expect a drop in US government default risk to increase the spread between corporate bond yields and Treasury yields. To examine this possibility, recall that Figure 5 plots the event study for the ICE BofA Option-Adjusted Spread index, and shows no statistically or economically significant increase in spreads. While this does not fully rule out that the effects we observe could be driven by changes in risk premia on US Treasuries, the risk premia on corporate bonds would need to move essentially in tandem.

A second explanation for our results could be that, while we are obtaining valid estimates for the impact of AI news at the dates we study, our event dates are in some sense non-representative. That is, it could be that the net effect of investor beliefs about AI has been to increase bond yields over the 2023-4 period, but that the particular event dates we've selected saw updates in the opposite direction. While we cannot rule this out, it is not clear to us why it would be the case: we include all dates from a well-defined universe (all major model release dates from a set of prominent AI firms), and it is not clear to us why the impact of information arriving at these dates should be directionally different, in aggregate, than that of AI information arriving at other dates in the same two year window.

A third possibility is that while we are accurately capturing market responses to AI news, the model in Section 5 implies a misleading interpretation of these results. There are a wide variety of reasons why reality may deviate from the fully-optimizing, complete market benchmark, including market incompleteness, a wide array of market frictions and constraints, behavioral deviations from rationality and optimization, and many more.²¹ To explain our results, an alternative story needs to explain economically large and apparently persistent yield changes in one of the deepest financial markets in the world. This suggests that alternative explanations could themselves be of considerable interest.

²¹For instance, perhaps asset managers face institutional risk-management constraints and expect a heightened level of volatility in the equity markets following AI model releases, leading to a shift of investments towards fixed-income instruments such as Treasury bonds. This explanation suggests a shift towards more liquid, short-maturity Treasuries, however, rather than the longer-maturity bonds where our effects are concentrated.

References

- Acemoglu, Daron and Pascual Restrepo**, “The wrong kind of AI? Artificial intelligence and the future of labour demand,” *Cambridge Journal of Regions, Economy and Society*, 2020, *13* (1), 25–35.
- **and Todd Lensman**, “Regulating transformative technologies,” *American Economic Review: Insights*, 2024, *6* (3), 359–376.
- , **David Autor, Jonathon Hazell, and Pascual Restrepo**, “Artificial intelligence and jobs: Evidence from online vacancies,” *Journal of Labor Economics*, 2022, *40* (S1), S293–S340.
- Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson**, “Artificial intelligence, firm growth, and product innovation,” *Journal of Financial Economics*, 2024, *151*, 103745.
- , — , **Alex X He, and James Hodson**, “Firm investments in artificial intelligence technologies and changes in workforce composition,” 2023, *31325*.
- Binsbergen, Jules Van, Wouter Hueskes, Ralph Koijen, and Evert Vrugt**, “Equity yields,” *Journal of Financial Economics*, 2013, *110* (3), 503–519.
- Board of Governors of the Federal Reserve System, US**, “Fitted Yield on 1-10 Year Zero Coupon Bond, Quoted on an Investment Basis, Inflation-Indexed [THREEFY1,THREEFY2,THREEFY3,THREEFY4,THREEFY5,THREEFY6,THREEFY7,THREEFY8,THREEFY9,THREEFY10],” FRED, Federal Reserve Bank of St. Louis 2025.
- , “Market Yield on U.S. Treasury Securities at 1,5,10,20,30-Year Constant Maturity, Quoted on an Investment Basis [DGS1, DGS5, DGS10, DGS20, DGS30],” FRED, Federal Reserve Bank of St. Louis 2025.
- , “Market Yield on U.S. Treasury Securities at 1,5,10,20,30-Year Constant Maturity, Quoted on an Investment Basis, Inflation-Indexed [DFII5, DFII10, DFII20, DFII30],” FRED, Federal Reserve Bank of St. Louis 2025.
- , “Nominal Broad U.S. Dollar Index [DTWEXBGS], retrieved from FRED, Federal Reserve Bank of St. Louis,” <https://fred.stlouisfed.org/series/DTWEXBGS> 2025.
- Borio, Claudio EV, Piti Disyatat, Mikael Juselius, and Phurichai Rungcharoenkitkul**, “Why so low for so long? A long-term view of real interest rates,” 2017.

- Bruce, E Ananth S and A Hansen**, “Uncovering the relationship between real interest rates and economic growth,” *Ann Arbor MI: University of Michigan Retirement Research Centre (MRRC)*, 2013.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson**, “Artificial intelligence and the modern productivity paradox,” *The economics of artificial intelligence: An agenda*, 2019, 23 (2019), 23–57.
- Cboe**, “Cboe Volatility Index (VIX),” 2025. 8/28/2025.
- Chiang, Wei-Lin, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li, Dacheng Li, Hao Zhang, Banghua Zhu, Michael Jordan, Joseph E. Gonzalez, and Ion Stoica**, “Chatbot Arena: An Open Platform for Evaluating LLMs by Human Preference,” 2024.
- Chow, Trevor, Basil Halperin, and J Zachary Mazlish**, “Transformative AI, existential risk, and real interest rates,” Technical Report, Working Paper 2024.
- Citigroup Global Markets**, “Citigroup Economic Surprise Index,” Available via Bloomberg Terminal (ticker: CESIUSD Index) 2025.
- Duffie, Darrell**, *Dynamic asset pricing theory*, Princeton University Press, 2010.
- Eisfeldt, Andrea L, Gregor Schubert, Miao Ben Zhang, and Bledi Taska**, “The labor impact of generative AI on firm values,” *Available at SSRN*, 2024, 4436627.
- Federal Reserve Bank of San Francisco**, “Daily News Sentiment Index,” <https://www.frbsf.org/research-and-insights/data-and-indicators/daily-news-sentiment-index/> 2023.
- Gil, Hamilton Galindo**, “Asset Pricing Impacts of AI in a Model with Heterogeneous Risk Aversion,” Technical Report, Working Paper 2024.
- Gürkaynak, Refet S, Brian Sack, and Jonathan H Wright**, “The TIPS yield curve and inflation compensation,” *American Economic Journal: Macroeconomics*, 2010, 2 (1), 70–92.
- Hamilton, James D, Ethan S Harris, Jan Hatzius, and Kenneth D West**, “The equilibrium real funds rate: Past, present and future,” Technical Report, National Bureau of Economic Research 2015.

- Hampole, Menaka, Dimitris Papanikolaou, Lawrence DW Schmidt, and Bryan Seegmiller**, “Artificial intelligence and the labor market,” Technical Report, National Bureau of Economic Research 2025.
- Ice Data Indices, LLC**, “ICE BofA 1-3, 3-5, 7-10, 10-15, 15+ Year US Corporate Index Effective Yield, [BAMLC1A0C13YEY, BAMLC2A0C35YEY, BAMLC3A0C57YEY, BAMLC4A0C710YEY, BAMLC7A0C1015YEY, BAMLC8A0C15PYEY],” FRED, Federal Reserve Bank of St. Louis 2025.
- , “ICE BofA 1-3, 3-5, 7-10, 10-15, 15+ Year US Corporate Index Option-Adjusted Spread, [BAMLC1A0C13YY, BAMLC2A0C35YY, BAMLC3A0C57YEY, BAMLC4A0C710YY, BAMLC7A0C1015YY, BAMLC8A0C15PYY],” FRED, Federal Reserve Bank of St. Louis 2025.
- Jackwerth, Jens Carsten and Mark Rubinstein**, “Recovering probability distributions from option prices,” *The journal of Finance*, 1996, 51 (5), 1611–1631.
- Jones, Charles I**, “The ai dilemma: Growth versus existential risk,” *American Economic Review: Insights*, 2024, 6 (4), 575–590.
- Kim, Don H. and Jonathan H. Wright**, “An Arbitrage-Free Three-Factor Term Structure Model and the Recent Behavior of Long-Term Yields and Distant-Horizon Forward Rates,” Technical Report, Working Paper 2005.
- Kokotajlo, Daniel, Scott Alexander, Thomas Larsen, Eli Lifland, and Romeo Dean**, “AI 2027,” 2025.
- Korinek, Anton and Donghyun Suh**, “Scenarios for the Transition to AGI,” Technical Report, National Bureau of Economic Research 2024.
- Lucca, David O and Emanuel Moench**, “The pre-FOMC announcement drift,” *The Journal of finance*, 2015, 70 (1), 329–371.
- Lunsford, Kurt G and Kenneth D West**, “Some evidence on secular drivers of US safe real rates,” *American Economic Journal: Macroeconomics*, 2019, 11 (4), 113–139.
- Mehra, Rajnish**, “Consumption-based asset pricing models,” *Annu. Rev. Financ. Econ.*, 2012, 4 (1), 385–409.

- Metaculus**, “When Will the First General AI Be Announced?,” <https://www.metaculus.com/questions/5121/> 2020.
- , “When Will Weakly General AI Arrive?,” <https://www.metaculus.com/questions/3479/date-weakly-general-ai-is-publicly-known/> 2020.
- Mollick, Ethan**, “Google’s Gemini Advanced: Tasting Notes and Implications,” 2024.
- Rogoff, Kenneth S, Barbara Rossi, and Paul Schmelzing**, “Long-run trends in long-maturity real rates, 1311–2022,” *American Economic Review*, 2024, *114* (8), 2271–2307.
- Savor, Pavel and Mungo Wilson**, “How much do investors care about macroeconomic risk? Evidence from scheduled economic announcements,” *Journal of Financial and Quantitative Analysis*, 2013, *48* (2), 343–375.
- Shapiro, Adam Hale, Moritz Sudhof, and Daniel J. Wilson**, “Measuring news sentiment,” *Journal of Econometrics*, 2022, *228* (2), 221–243.
- Trammell, Philip and Anton Korinek**, “Economic growth under transformative AI,” Technical Report, National Bureau of Economic Research 2023.
- van Binsbergen, Jules, Michael Brandt, and Ralph Koijen**, “On the timing and pricing of dividends,” *American Economic Review*, 2012, *102* (4), 1596–1618.
- Webb, Michael**, “The impact of artificial intelligence on the labor market,” *Available at SSRN 3482150*, 2019.
- Wolfers, Justin and Eric Zitzewitz**, “Prediction markets,” *Journal of economic perspectives*, 2004, *18* (2), 107–126.

Appendix

A Exchange Rate Responses

Figure 8 shows an event study for a broad trade-weighted US dollar index from FRED (Board of Governors of the Federal Reserve System, US, 2025d). As this plot shows, on median the model releases in our sample saw a weakening of the dollar, consistent with lower demand for the dollar following the fall in interest rates estimated in the main text. These declines are significant relative to the placebo distribution, and again persist through 15 days after the model release.




Figure 8: Median change in trade-weighted US Dollar index (relative to fifteen days before event). Median taken across AI release events in the 2023 and 2024 calendar years. Placebo distribution recomputes statistics on dates drawn uniformly at random from sample period.


B Robustness Checks

B.1 Dropping Events

To examine robustness of our results to small changes in our date series, Figures 9 and 10 show results obtained by dropping all combinations of one or two (in the top and bottom panels of Figure 9, respectively) or three (in Figure 10) event dates from our series. The results are directionally similar to those based on the full sample, particularly for longer-maturity bonds, even when dropping two or three event dates from the sample. Statistical significance is more fragile, but is at most horizons robust to dropping a single event and, at some horizons and for maturities longer than 10 years, robust to dropping any combination of two or three events.

Drop-1 Robustness: Median change across events

15 drop-1 series (green lines); Full sample (blue line); Placebo bands: 99%, 95%, 90% (grey dashed); Dark grey = placebo mean



Drop-2 Robustness: Median change across events

105 drop-2 series (green lines); Full sample (blue line); Placebo bands: 99%, 95%, 90% (grey dashed); Dark grey = placebo mean




Figure 9: Median change in yields (relative to fifteen days before event) for US treasuries. Blue line takes median across AI release events in the 2023 and 2024 calendar years, while green lines correspond to alternative series obtained by dropping one (in top panel) or two (in bottom panel) event dates from series. Placebo distribution recomputes statistics on dates drawn uniformly at random from sample period.




Figure 10: Median change in yields (relative to fifteen days before event) for US treasuries. Blue line takes median across AI release events in the 2023 and 2024 calendar years, while green lines correspond to alternative series obtained by dropping three event dates from series. Placebo distribution recomputes statistics on dates drawn uniformly at random from sample period.

Table 5: Major Tech Conference Dates

Date	Company	Conference
<i>2023 Conferences</i>		
3/1/2023	Tesla	Tesla AI Day
3/23/2023	Nvidia	Nvidia GTC
5/10/2023	Google	Google I/O
5/25/2023	Microsoft	Microsoft Build
6/9/2023	Apple	Worldwide Developer Conference
9/28/2023	Meta	Meta Connect
12/1/2023	Amazon	AWS re:Invent
<i>2024 Conferences</i>		
3/21/2024	Nvidia	Nvidia GTC
4/22/2024	Tesla	Tesla AI Day
5/14/2024	Google	Google I/O
5/23/2024	Microsoft	Microsoft Build
6/14/2024	Apple	Worldwide Developer Conference
9/26/2024	Meta	Meta Connect
12/6/2024	Amazon	AWS re:Invent

Notes: This table presents the major annual tech conferences from the “Magnificent Seven” companies used in our event study analysis. For multi-day conferences, we use the final day.

B.2 Alternative Placebos

In this appendix, we explore the extent to which other date series could potentially explain the results we find around AI model release dates by employing these alternative series as placebo dates. Specifically, we consider date series corresponding to (i) FOMC meetings (ii) major conferences series held by the so-called “Magnificent Seven” major tech firms (listed in Table 5) (iii) earnings release dates for the same seven firms, (iv) Consumer Price Index release dates, (v) Bureau of Labor Statistics Employment Situation Series release dates (vi) US Census Bureau Advance Monthly Sales for Retail and Food Services release dates and (vii) US Treasury auctions for 10, 20, or 30 year bonds.

For the first two series, once we limit attention to our analysis window of calendar years 2023 and 2024, we have no more events than we do AI model releases. Hence, we lack sufficient dates to generate placebo distributions, and simply plot (in Figure 11) the median change in yields, based on these alternative date series, against the median change in yields across the model release dates used in our analysis. Across both series we see that these events are not associated with a decline in long-maturity Treasury yields of the sort observed around the model release dates.

For the remaining placebo date series we have sufficiently many events to produce non-trivial placebo distributions. There is considerable clustering in earnings announcements, so given our long event windows we group earnings announcements within a single calendar week into a single event, with the event date recorded as the final earnings announcement in that week. For each placebo series we draw dates at random (without replacement) from the corresponding date lists, and obtain the placebo distributions shown in Figures 12-14. Once again, we find that behavior of US Treasury yields around AI model release dates, particularly for long-maturity bonds, appears statistically and economically different from behavior around these alternative date series.

It is worth emphasizing that statistical insignificance of our AI events relative to these alternative placebo series would *not* suffice to imply that these alternative event series explain our results, since these alternative event dates would also need to be sufficiently correlated with our events. On the other hand, statistical significance for our results relative to these alternative placebos implies that even if our event dates were a subset of these alternative dates (which they are not, for any of these series), the alternative series would not explain our results.





Figure 11: Median change in yields (relative to fifteen days before event) for US treasuries. Median taken across AI release events in the 2023 and 2024 calendar years. First panel placebo takes median over FOMC meetings in the 2023 and 2024 calendar years. Second panel placebo takes median over “Magnificent Seven” tech company conference dates.

Median change across events - Tech Earnings Dates

Placebo bands: 99%, 95%, 90% (draws = 5000); Dark grey line = placebo mean - Placebo from tech dates only



Median change across events - CPI Dates

Placebo bands: 99%, 95%, 90% (draws = 5000); Dark grey line = placebo mean - Placebo from CPI only





Figure 12: Median change in yields (relative to fifteen days before event) for US Treasuries. Median taken across AI release events in the 2023 and 2024 calendar years. First panel placebo takes median over “Magnificent Seven” earnings announcement dates. Second panel placebo takes median over CPI release dates.

Median change across events - Jobs Report Dates

Placebo bands: 99%, 95%, 90% (draws = 5000); Dark grey line = placebo mean - Placebo from Jobs only



Median change across events - Retail Sales Dates

Placebo bands: 99%, 95%, 90% (draws = 5000); Dark grey line = placebo mean - Placebo from RetailSales only




Figure 13: Median change in yields (relative to fifteen days before event) for US Treasuries. Median taken across AI release events in the 2023 and 2024 calendar years. First panel placebo takes median over BLS employment situation report release dates. Second panel placebo takes median over retail sales report release dates.




Figure 14: Median change in yields (relative to fifteen days before event) for US Treasuries. Median taken across AI release events in the 2023 and 2024 calendar years. Placebo takes median over auction dates for 10, 20, and 30 year US Treasuries.

B.3 Controlling for Other News

As discussed in the main text, in order to control for non-AI news which arrived during our sample period, in this appendix we residualize changes in bond yields against (i) the Citigroup US Economic Surprise Index (Citigroup Global Markets, 2025) (ii) the Cboe’s VIX volatility index (Cboe, 2025), and (iii) the Federal Reserve Bank of San Francisco Daily News Sentiment Index (Shapiro et al., 2022; Federal Reserve Bank of San Francisco, 2023). For each control series and each Treasury series, we regress changes in bond yields against a constant, the control series, and 15 lags of the control series, where we run this regression of the full 2023-4 analysis window (using control observations from 2022 to avoid dropping yield observations at the start of the sample). This removes the component of yield changes which is linearly predicted by the current and lagged control series and so, to the extent the control series captures non-AI news which may affect bond yields, removes this from the yield series. We then re-run our placebo test using the residualized yield series, now testing the null hypothesis of no effect of AI model releases on the residualized yield series.

Each of our control series has advantages and disadvantages. The US Economic Surprise Index captures the deviation of economic data releases from forecasters’ predictions, and thus is a convenient summary for the unexpected component of official data releases. While the series can be downloaded from Bloomberg, however, the precise construction of the series is proprietary. By contrast, the construction of the VIX index from options data is more transparent, but the precise economic factors which drive option-implied volatility are less clear. In particular, it may be that AI news impacts implied volatility so the VIX is a so-called “bad control.” Similarly, the construction of the Daily News Sentiment Index is based on public methodology (Shapiro et al., 2022), but economic sentiment may be impacted by AI releases.

Figures 15 and 16 show event study plots based on our US Treasury series where we control for each series separately (in Figure 15 and the first panel of Figure 16) or for all three series at once (in the second panel of Figure 16). As we see from these figures, our results are broadly similar across all sets of controls, with economically large and, at least at some horizons, statistically significant declines in bond yields following AI model releases, particularly for longer-maturity bonds. That said, the level of statistical significance does vary across specifications, and is lowest when we control for all three series simultaneously.




Figure 15: Median change in yields (relative to fifteen days before event) for US Treasuries, residualized against Citigroup US Economic Surprise Index (in first panel) and VIX option-implied volatility index (in second panel). Median taken across AI release events in the 2023 and 2024 calendar years. Placebo distribution recomputes statistics on dates drawn uniformly at random from sample period.




Figure 16: Median change in yields (relative to fifteen days before event) for US Treasuries, residualized against Federal Reserve Bank of San Francisco Daily News Sentiment Index (in first panel) and all three control series (in second panel). Median taken across AI release events in the 2023 and 2024 calendar years. Placebo distribution recomputes statistics on dates drawn uniformly at random from sample period.

Table 6: AI Model Release Dates

Date	AI Laboratory	Model
<i>2022 Releases</i>		
11/30/2022	OpenAI	ChatGPT 3.5
<i>2025 Releases</i>		
01/20/2025	DeepSeek	DeepSeek R1
02/19/2025	xAI	Grok 3
02/24/2025	Anthropic	Claude 3.7
02/28/2025	OpenAI	ChatGPT 4.5
03/25/2025	Google	Gemini 2.5
04/16/2025	OpenAI	o3
05/22/2025	Anthropic	Claude 4

Notes: This table presents the additional AI model releases used in our extended event study analysis, covering from October 2022 through May 2025.

B.4 Alternative Analysis Windows

To complement the results for the 2023-2024 period reported in the main text, this appendix separately reports results for calendar years 2023 and 2024, as well as an extended analysis period from October 2022 through May 2025.

The top panel of Figure 17 reports results for 2023, while the bottom panel reports results for 2024. As Figure 17 shows, our analysis applied in calendar year 2023 finds statistically significant effects with even larger magnitudes than our pooled analysis. By contrast, in 2024 we find results which, while directionally similar to our full-sample results, and similar in magnitude at many horizons, are statistically insignificant at conventional levels by the end of our analysis window.

Figure 18 extends our analysis window to cover the period from October 2022 through May 2025, including the additional model releases detailed in Table 6, and presents median changes in treasury yields. We see that our results for this extended period are similar to those for our baseline analysis.




Figure 17: Median change in yields (relative to fifteen days before event) US Treasuries. Top panel takes median across six AI model release dates in 2023, while bottom panel takes median across nine model release dates in 2024. Placebo distribution in each panel recomputes statistics on dates drawn uniformly at random from the same year.




Figure 18: Median change in yields (relative to fifteen days before event) US Treasuries. Median taken across AI release events from October 2022 through May 2025. Placebo distribution recomputes statistics on dates drawn uniformly at random from sample period.

B.5 Means vs. Medians

In the main text we focus on results for median and median absolute changes in yields around major AI model releases. To complement those results, Figures 19-21 report a versions of Figures 2-4 instead using the mean change

$$\text{MeanChange}_s = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \Delta y_{t,s}$$

and the mean absolute change

$$\text{MeanAbsChange}_s = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} |\Delta y_{t,s}|.$$

The findings are qualitatively similar to those based on medians, with strong declines in yields, particularly at longer maturities.




Figure 19: Mean and mean absolute change in yields (relative to fifteen days before event) for constant-maturity US Treasury Bonds. Mean taken across AI release events in the 2023 and 2024 calendar years. Placebo distribution recomputes statistics on dates drawn uniformly at random from sample period.




Figure 20: Mean and mean absolute change in yields (relative to fifteen days before event) for constant-maturity TIPS. Mean taken across AI release events in the 2023 and 2024 calendar years. Placebo distribution recomputes statistics on dates drawn uniformly at random from sample period.




Figure 21: Mean and mean absolute change in yields (relative to fifteen days before event) for corporate bond indices. Mean taken across AI release events in the 2023 and 2024 calendar years. Placebo distribution recomputes statistics on dates drawn uniformly at random from sample period.

C Additional Results For Section 5

C.1 Additional Derivations

Note that by the law of iterated expectations and our conditional independence assumptions

$$\begin{aligned}\mathbb{E}_t[M_{t,t+h}] &= \mathbb{P}_t(t+k \leq T) \mathbb{E}_t[M_{t,t+h} | t+k \leq T] = \\ & \mathbb{P}_t(t+k \leq T) \mathbb{E}_t \left[\left(\frac{C_{t+k}}{C_t} \right)^{-\gamma} \beta^h (1+g)^{-(h-k)\gamma} \left(\prod_{s=k+1}^h X_{t+s} \right)^{-\gamma} 1_{\{t+h \leq T\}} \middle| t+k \leq T \right] = \\ & \mathbb{P}_t(t+k \leq T) \beta^h \mathbb{E}_t \left[\left(\frac{C_{t+k}}{C_t} \right)^{-\gamma} \left(\prod_{s=k+1}^h X_{t+s} \right)^{-\gamma} \middle| t+k \leq T \right] \times \\ & \mathbb{E}_t \left[(1+g)^{-(h-k)\gamma} \middle| t+k \leq T \right] \mathbb{P}_t(t+h \leq T | t+k \leq T).\end{aligned}$$

By the formula for the mean of a lognormal, however,

$$\mathbb{E}_t \left[(1+g)^{-(h-k)\gamma} \middle| t+k \leq T \right] = \exp \left(-\gamma(h-k)\mu_t + \frac{\gamma^2(h-k)^2\sigma_t^2}{2} \right),$$

so we can re-write the expected h -period-ahead SDF as

$$\begin{aligned}\mathbb{E}_t[M_{t,t+h}] &= \\ & \beta^h (1-\delta_t)^{(h-k)} \exp \left(-\gamma(h-k)\mu_t + \frac{\gamma^2(h-k)^2\sigma_t^2}{2} \right) \mathbb{P}_t(t+k \leq T) \mathbb{E}_t \left[\left(\frac{C_{t+k}}{C_t} \right)^{-\gamma} \left(\prod_{s=k+1}^h X_{t+s} \right)^{-\gamma} \right].\end{aligned}$$

Consequently, if we consider the difference of log expected SDFs, $\log(\mathbb{E}_t[M_{t,t+h}]) - \log(\mathbb{E}_t[M_{t,t+k}])$ we get

$$\begin{aligned}(h-k)\log(\beta) + (h-k)\log(1-\delta_t) - \gamma(h-k)\mu_t + \frac{\gamma^2}{2}(h-k)^2\sigma_t^2 \\ + \log \left(\frac{E_t \left[\left(\frac{C_{t+k}}{C_t} \right)^{-\gamma} \left(\prod_{s=k+1}^h X_{t+s} \right)^{-\gamma} \right]}{E_t \left[\left(\frac{C_{t+k}}{C_t} \right)^{-\gamma} \right]} \right).\end{aligned}$$

Note, however, that since $y_{t,t+h} = \mathbb{E}_t[M_{t,t+h}]^{-\frac{1}{h}}$, we have

$$\log(f_{t+k,t+h}) = \frac{\log(\mathbb{E}_t[M_{t,t+k}]) - \log(\mathbb{E}_t[M_{t,t+h}])}{h-k} =$$

$$-\log(\beta) - \log(1 - \delta_t) + \gamma\mu_t - \frac{\gamma^2}{2}(h-k)\sigma_t^2 - \frac{1}{h-k} \log \left(E_t \left[\frac{\left(\frac{C_{t+k}}{C_t}\right)^{-\gamma}}{\left[\frac{C_{t+k}}{C_t}\right]^{-\gamma}} \left(\prod_{s=k+1}^h X_{t+s} \right)^{-\gamma} \right] \right),$$

as stated in the main text.

C.2 Diagnostic Plot for Choice of k

As discussed in Section 5, under our simplified model the difference in differences in one period-ahead log forward yields, $\text{DID}(\log(f_{t+h,t+h+1}); \mathcal{T}, \mathcal{A})$ for $h \in \{0, \dots, 9\}$, should be approximately linear in h for $h \geq k$, which we use to guide our choice of k . Figure 22 plots this one-period-ahead difference in difference, and shows that for $h \geq 4$ linearity appears a good approximation.




Figure 22: Difference in differences for one-period ahead log forward rate starting at period h , $DID(\log(f_{t+h,t+h+1}); \mathcal{T}, \mathcal{A})$. Placebo distribution recomputes statistics on dates drawn uniformly at random from sample period.